A Learning Algorithm for Change Impact Prediction: Experimentation on 7 Java Applications

Vincenzo Musco, Antonin Carette, Martin Monperrus, Philippe Preux
University of Lille and INRIA
Lille, France
Contact: vincenzo.musco@inria.fr

arXiv:1512.07435v1 [cs.SE] 23 Dec 2015

Abstract

Change impact analysis consists in predicting the impact of a code change in a software application. In this paper, we take a learning perspective on change impact analysis and consider the problem formulated as follows. The artifacts that are considered are methods of object-oriented software; the change under study is a change in the code of the method, the impact is the test methods that fail because of the change that has been performed. We propose an algorithm, called LCIP that learns from past impacts to predict future impacts. To evaluate our system, we consider 7 Java software applications totaling 214,000+ lines of code. We simulate 17574 changes and their actual impact through code mutations, as done in mutation testing. We find that LCIP can predict the impact with a precision of 69%, a recall of 79%, corresponding to a F-Score of 55%.

1 Introduction

Change impact analysis consists in predicting the impact of a code change in a software application [1, 2]. There are different kinds of impacts: the software modules that may be broken [1], the documentation that must be impacted [3], the developers that must keep informed about this change [4]. In this paper we focus on the first kind: impacts at the code level.

There are two important challenges for change impact analysis. First, it must scale to the size of today's software systems, that often consists of thousands of modules. Second, the predictions must be accurate. The accuracy of impact prediction has two dimensions: whether predicted impacted elements are actually impacted (false positives) and whether actually impacted elements are all predicted (false negatives). This boils down to the classical information retrieval measures: precision, recall, and F-score. To our knowledge, it is an open research question to design a change impact prediction system that addresses both challenges: scaling to large systems while retaining a high accuracy (high precision, high recall, high F-score).

Our intuition is to take a learning perspective on the problem of change impact analysis. By learning, we mean that the change impact analysis could exploit previous actual impacts to form some kind of knowledge; this knowledge then guides the prediction of the future impacts. To this extent, a learning-based change impact analysis technique is the opposite of *a priori* approaches which design upfront the whens and hows the impact propagates, such as those based on program dependence graphs [5] to only name a few. To our knowledge, we are the first to take a learning viewpoint on the problem.

In this paper, we consider change impact analysis formulated as follows. The artifacts that are considered are methods of object-oriented software; the change under study is a change in the code of the method, the impact is the test methods that fail because of the change that has been performed. We propose an algorithm, called LCIP that learns from past impacts to predict future impacts. For instance, the algorithm may learn that a change in method `foo` is always followed by failure observation in `bar`. Our algorithm is based on the call graph, where each node is a method and each edge a potential call between two methods, as observed statically. We consider a standard Class-Hierarchy-Analysis (CHA) call graph. The learning strategy consists of decorating each call graph edge with a weight between zero and one. This weight represents the likelihood of an edge to propagate an impact. The weight is updated based on actual impacts that are given as training data to the system. To our knowledge, our problem statement does not correspond to any standard problem in the machine learning literature.

To evaluate our system, we consider 7 Java software applications totaling 214,000+ lines of code. We simulate 17574 changes and their actual impact through code mutations, as done in mutation testing. With ten-fold cross validation, we measure the precision, recall and F-Score of change impact prediction. We compare our system to two algorithms: one being a standard transitive closure on the call graph, the other one is a basic learning strategy that we introduce for the sake of understanding the actual learning that happens in LCIP. We find that LCIP can predict the impact with a precision of

69%, a recall of 79%, corresponding to a F-Score of 55%. This validates our intuition that learning can be done for change impact analysis. In addition, our approach does not trade performance for accuracy: learning lasts in average 26 seconds and prediction 256 milliseconds. This indicates that our approach may scale-up to really-large systems.

To sum up, our contributions are:

- an algorithm for learning the impact of software changes. It is based on decorating call graph edges with a weight representing a likelihood to propagate a change.
- An experimentation made on 7 popular Java programs, totaling 214,000+ lines of codes. The experiment shows that a learning approach is a promising approach for change impact prediction.
- We publish all our code as open-source. The frameworks for mutation, impact prediction and learning can foster future research in this novel research direction.

The remainder of this paper is structured as follows. In Section 2 we present concepts used in this paper and introduce our learning approach. In Section 3 we present our framework and dataset. We also pose our research questions and answer them by presenting and analyzing our experimental results. In Section 4, we present works related to our paper in fields such as machine learning, impact analysis and mutation testing. In Section 5 we conclude this paper.

2 Contribution

We present a new learning algorithm for change impact analysis. This algorithm is based on decorating call graphs with weights on the edges.

2.1 Concepts

2.1.1 Change Impact Analysis

Software is made of interconnected pieces of code (*e.g.* methods, classes, *etc.*). Through those connections, the effects of a change to a given part of the code can propagate to many other parts of the software. Those other parts, which can be potentially anywhere in the software, can then be impacted by the initial change, acting like a ripple. Change Impact Analysis (*a.k.a.* CIA) is defined as "the determination of potential effects to a subject system resulting from a proposed software change" [6].

In this paper, we use Bohner's definition of the basic software change impact analysis process [6]. Assuming a change has been performed, Bohner defines the following sets used in impact analysis: (i) the "starting impact set" (SIS) is the list of software parts which can be impacted by the change; (ii) the "candidate impact set" (CIS) (also called the "estimated impact set" [7]) is the list of software parts predicted as impacted by a change impact analysis technique; (iii) the "actual impacted set" (AIS) is the list of parts of the software which are actually impacted by the change; (iv) the "discovered impact set" (DIS) is the list of missed impacts by the technique (*i.e.* the false negatives); (v) the "false positive impact set" ($FPIS$) is the list of over-estimated impact returned by the technique (*i.e.* false positives).

2.1.2 Mutation Analysis

Mutation analysis [8] consists of assessing test suite quality by applying minor changes to a software resulting in slightly modified versions of the software which are called *software mutants*. These changes result from the application of a *mutation operator* which describes what should be changed in the code, and how it should be modified.

A standard mutation testing scenario is that 1) the software tests are run on software mutants to ensure the change is noticed by at least one test and 2) If no test kills a mutant, complementary tests are written.

2.1.3 Call graph

Grove et al. [9] define "*the program call graph [as] a directed graph that represents the calling relationships between the program's procedures (...) each procedure is represented by a single node in the graph. Each node has an indexed set of call sites, and each call site is the source of zero or more edges to other nodes, representing possible callees of that site*". This is the definition taken in this paper.

As we want to study Java applications which are object oriented, we also take into consideration the class hierarchy analysis (*a.k.a.* CHA), *i.e.* the inheritance and usage of interfaces. This means we add nodes/edges for based on polymorphism.

2.2 Approach Overview

In this paper, we use mutation injection to artificially introduce a change in the software. We then consider the tests that are impacted by a mutation.

We produce a list S of mutant programs s' which are new versions of a program s which differs in a very simple way from s. Thus, running the tests on the mutant software will give us the actual impact AIS set which corresponds to the failing tests.

We use a call graph such as defined in Section 2.1.3 to do change impact analysis. Given a call graph and a changed node (*i.e.* a node in which a change has occurred), an impact prediction algorithm computes the candidate impacted set CIS composed of all test nodes which may be broken by the change.

Figure 1 illustrates a basic example of a call graph. Three types of nodes are presented: application

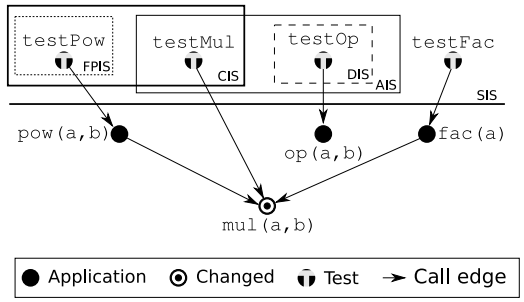

Figure 1: Example of a call graph in which a change has been introduced. The class graph includes application nodes, test nodes and call edges. The rectangles illustrate Bohner's sets.

nodes (plain circle), test nodes (circle with a T) and the changed node (double circle). We have four methods: `mul` which computes a multiplication and two methods depending on it: `pow` and `fac` which computes respectively the power and the factorial. Then, a test function is defined for each method. The `op` method uses reflection mechanism to call `mul`, which is not resolved statically, resulting in the absence of edge between `op` and `mul`.

Consider that an error is injected in the `mul` method. Bohner's sets are illustrated in the figure. The SIS set is made of all software test nodes. The top part of the horizontal line shows the SIS set. The AIS set (thin rectangle) is composed of the test nodes which fail when running on the mutant. The CIS set (thick rectangle) is composed of the candidate test nodes determined by a virtual change impact prediction technique (in this example it is an arbitrary set). The DIS set (dashed rectangle) shows a false negative test due to the absence of edge in the call graph. The $FPIS$ set (dotted rectangle) shows a false positive case: the `pow` method seems to be impacted according to the change impact analysis technique (as it is part of the CIS set), but when running the tests, it did not fail.

2.3 Learning for Impact Prediction

We use the call graph to estimate the change impact set. However, error propagation through an edge is not a mere black-or-white process: some edges may propagate all errors, other edges may never propagate errors. There may also exist edges that propagate errors but only sometimes, in particular cases. Thus, we propose a stochastic call graph in which we use weights on the call graph edges. These weights range from 0 to 1 where 0 means that the edge never propagates the change, while 1 means that the edge always propagates it. Any value in between means the propagation sometimes occurs and sometimes not. The initial weights are set to 0 meaning that we start by assuming that no impact is propagated at all.

In this paper, we propose the Learning Change Impact Prediction (LCIP) algorithm. It is composed of a top-level learning loop and an algorithm for updating edge weights called Algorithm Dichotomic. The top-level impact prediction algorithm is shown in Algorithm 1. This algorithm takes as input the node n corresponding to the point in the code where the change has occurred and a threshold value th lying in the range $[0, 1]$. It returns a candidate impacted set CIS composed of all tests predicted as impacted. To do so, starting at the node being changed, the graph edges are followed to determine which tests can be reached. The weights are used to prune some edges which are unlikely to propagate the change, according to the threshold value (line 6). If the weight is lower than the threshold value, the error does not propagate across the edge; otherwise the edge propagates the error.

Algorithm 1: The impact prediction algorithm that uses the learned weights of the call graph

Input: n the node belonging to the call graph , th the treshold

```
1  begin
2      CIS ← {}
3      for each node i connecting to n do
4          if i is not visited then
5              mark node i as visited
6              if w_i ¿= th then
7                  if i is a test node then
8                      CIS ← CIS ∪ {i}
9                  CIS ← CIS ∪ visit(i)
10     return CIS
```

In order to learn the best weights for the call graph edges, we use information gathered from runs of the program with various changes: we inject a change in the program at some mutation point, and we observe at which test points the change has propagated. This provides us with a set of "examples" that a machine learning algorithm uses as its input. The algorithm estimates the weights to assign to the edges of the call graph in order to predict accurately the propagation of changes. Thus, our algorithm requires as input data: (i) a call graph G such as described in 2.1.3; (ii) a set M made of tuples containing all mutation points and impacted tests such as:

$$M = \{\{m, AIS\} \mid \forall S\} \quad (1)$$

where S is the set of mutated versions of the software, m is the method on which the change has occurred and AIS is a set of test methods which have been impacted by the change on run time (by running test cases).

All elements contained in the set M forms a dataset. We split this dataset in two parts known as the training set and the testing set; the training

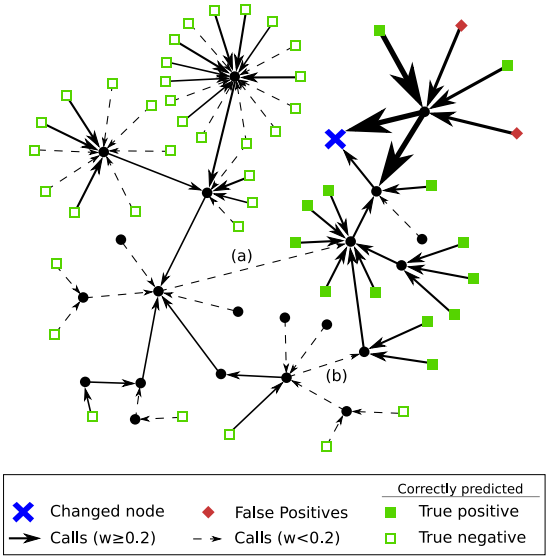

Figure 2: An illustration of impact prediction based on weighted call graphs. Edges with a low weight (< 0.2) are considered as non-propagating the impact of the change.

set is used to learn the weights. The testing set is used to assess the weights that have been learned.

Our top-level learning impact algorithm is shown on Algorithm 2. For all mutants occurring at mutation point m (line 2) and each test case t (line 3), we update the weight of edges belonging to all paths from m to t following an update algorithm (line 4). In this paper, we propose two algorithms for updating the weights, that we now present.

Algorithm 2: The top-level impact learning algorithm using call graph. *update_weight* is a sub-algorithm which updates the weights.

Input: G the call graph and M the data defined in Section 2.3

```
1 begin
2     for each {m, AIS} ∈ M do
3         for each t ∈ AIS do
4             update_weight(m, t, G)
```

Figure 2 illustrates our impact analysis technique based on decorating call graphs with weights. It is based on real data obtained in our experiments. Assume that a change has occurred at the method denoted by a blue cross. The edges are method calls, and the thickness of the edges represents the weight of the edge after learning. A thicker edge means a weight close to 1 and a thinner means a weight close to 0.2. The dashed edges are those which have a weight smaller than 0.2 and which are not considered for propagation. The black nodes are application nodes. The green squares and red diamonds represent tests predicted as impacted by our approach. The prediction is of high accuracy because the weights of the two edges (a) and (b) is low after learning. Consequently, the impact of the change is stopped and does not ripple to the left-hand side of the graph. On the contrary, a basic impact prediction based on a transitive closure on the call graph would predict far too much tests (all the true negatives would become false positives).

2.3.1 Binary update algorithm

This algorithm assigns either 0 or 1 to the edge weights as follows. If at least one impact has been observed between the mutation point and the test during test case execution, then edges belonging to all paths going from the former to the latter are labeled as 1, otherwise, these edges are labeled with a 0. Thus, this apporach considers that if an edge has once propagated an error, it will always do so (as 1 is always greater or equal to the threshold, no matter its value). Algorithm 3 formalizes this idea. This algorithm is deterministic. Thus, running the same algorithm on the same data several times will always produce the same results.

Algorithm 3: Algorithm Binary for updating the edge weights in paths between the mutation point and the test methods.

Input: m the method on which mutation has occurred, t the considered impacted test and G the call graph

```
1 begin
2     for each edge in all paths from m to t in G
        do
3         w_edge ← 1
```

2.3.2 Dichotomic update algorithm

This algorithm updates the weights according to an estimation of the probability that a test would be broken by a change. This estimation is based on training data as follows:

$$p_{t,m} = \frac{\alpha_t}{\beta_m}; \tag{2}$$

where, α_t is the number of times the test t is impacted for all mutations occurring on the same point m and β_m is the number of times the method is mutated.

The idea of Algorithm Dichotomic is to slowly converge to $p_{t,m}$ example after example. Algorithm 4 is the corresponding algorithm. For each training example, the weight w of each edge which belongs to a path between a mutation node m and an impacted test node t is computed in a dichotomic way: the new weight is the mean value between the current weight and the empirical probability.

Table 1: Descriptive statistics about the 7 projects considered in this paper.

				Call graph					
				without CHA			with CHA		
Project	Version	Commit	LOC	#nodes	#edges	#paths	#nodes	#edges	#paths
Codec	1.11	r1676715 (svn)	17531	1373	2242	2227	1410	2360	3708
Collections	4.1	r1610049 (svn)	55081	5570	9335	1393	6271	12130	2260
IO	2.5	r1684201 (svn)	26528	2416	4059	1163	2494	4283	2047
Lang	3.5	#6965455 (git)	67509	6105	9312	26236	6195	9653	36412
Shindig	2.5.3	r1687149 (svn)	15710	1570	2465	999	1682	2713	2096
SonarQube	5.2	#1385dd3 (git)	29249	3301	5792	1075	3343	5907	1258
Spojo	1.0.7	#8fb2194 (git)	3371	330	890	1160	341	923	1350
Total			**214979**	**20665**	**34095**	**34253**	**21736**	**37969**	**49131**

Algorithm 4: Algorithm Dichotomic for updating the edge weights between the mutation point and the test methods.

Input: m the method on which mutation has occurred, t the considered impacted test, $p_{t,m}$ the empirical probability for (m, t) and G the call graph

```
1 begin
2     for each edge in all paths from m to t in G do
3         w_edge ← (w_edge + p_{t,m})/2
```

3 Evaluation

We now present the evaluation of our novel learning-based technique for impact prediction.

3.1 Research Questions

We are especially interested in answering the following research questions.

Research Question 1 *Do our learning algorithms improve the accuracy of impact prediction?* In this question, we want to determine whether our prediction algorithm based on learning can improve the prediction scores compared to the use of a baseline transitive closure prediction technique described in 3.2.2.

Research Question 2 *For the dichotomic approach, is the threshold value used for prediction important?* As presented in Section 2.3, the impact prediction algorithm requires a threshold value to tell whether an edge propagate an impact or not. This research question determines the importance of this value for the dichotomic and enables us to set the best threshold value.

Research Question 3 *If we do not take into consideration the class hierarchy analysis in the call graph, what is the impact on the learning time and the prediction accuracy of dichotomic algorithm?* In this paper we consider call graphs with class hierarchy analysis. However, the class hierarchy analysis can introduce a great amount of additional edges and paths as seen in Table 1. This additional data may have an impact on learning time. In this research question, we want to determine if we can use a lighter version of the call graph (by removing the inheritance and interfaces) instead of the full call graph. Answering to this question will help us to assess whether our approach can scale to large projects.

3.2 Experimental protocol

We explain how we evaluate our approach, the dataset and the configuration we use.

The evaluation follows several steps: (i) we create mutants for a Java software; (ii) we extract the corresponding call graphs with and without class hierarchy analysis; (iii) we split the dataset of mutants in a training set and a testing set. (iv) we run our learning algorithms based on the mutants of the training set. This results in weights. (v) we predict the impact set for each mutant. (vi) we compute performance metrics by comparing the predicted impact set and the actual impact set.

In addition, we use 10-fold cross-validation [10]: for each software, we partition the mutants into 10 subsets of equal size. We take 9 subsets to train the model (with Algorithm 2) and the one remaining is used to assess the model (with Algorithm 1). This process is run 10 times with different random ordering of the mutants at each run. We compute the mean value of the evaluation metrics considered over these 10 runs.

Our mutation, learning and evaluation framework is open-source and freely available online[1].

3.2.1 Evaluation Metrics

To assess and compare the performance of our impact prediction techniques, we use metrics computed on CIS and AIS presented by Arnold and Bohner [7]. For instance the size of the CIS has to be as close as possible to the size of the AIS as it quantifies the amount of elements retrieved by the impact

[1] `https://github.com/v-m/ChangePropagation`, `https://github.com/k0pernicus/PropL`

analysis. For the sake of performance evaluation, we can express it in terms of precision, recall and F-score [11]. There is one precision, recall and F-score per mutant of the testing set.

The *precision* P is the proportion of predicted tests which are actually impacted; P is computed using Equation (3). The *recall* R is the proportion of truly impacted tests that are retrieved; R is computed using Equation (4).

$$P = \frac{|AIS \cap CIS|}{|CIS|}, \tag{3}$$

$$R = \frac{|AIS \cap CIS|}{|AIS|}. \tag{4}$$

To make comparison easier, the *f-score* F combines both metrics P and R by computing their harmonic mean as shown in Equation (5).

$$F = 2 \times \frac{P \times R}{P + R}; \tag{5}$$

Since precision, recall and f-score are computed for each mutation, it is necessary to consider a mean value over all mutants of the testing set. *Our key goal is to improvement the f-score of impact prediction as it takes into consideration both precision and recall.*

3.2.2 Baseline

As a baseline for change impact analysis, we compute the basic call graph transitive closure obtained from the changed node [1]. The result of the transitive closure will be a list of all nodes potentially impacted by the change. Since the impact is only computed on test nodes, we remove application nodes from the impact nodes.

3.2.3 Dataset and Setup

In this paper we have selected 7 *Java* software packages: *Apache Commons Lang*, *Apache Commons Collections*, *Apache Commons Codec*, *Apache Commons Io*, *Apache Shindig*, *Spojo*, *SonarQube*. When the project is made of several sub-projects, we consider only the main one. Table 1 reports the key descriptive statistics about these projects:

- the first, second, and third columns give the name, the version, and the commit-id or the revision number (depending on whether the project is hosted on a git or subversion repository) of the software being analyzed;
- the fourth column indicates the number of lines of code of the project (computed using `cloc` [2]);
- the fifth, sixth, and seventh columns respectively give the number of nodes, the number of edges, and the number of paths between all mutation points and all tests in the call graphs when not considering class hierarchy analysis;
- the eight, ninth, and tenth columns give the same metrics for the call graphs using class hierarchy analysis.

3.2.4 Mutation Operators

We consider 5 mutation operators which are those presented by King and Offutt [12] adapted to the Java language (as those operators are firstly intended for Fortran). As shown by Offutt et al., these five mutation operators are sufficient to effectively implement mutation testing [13]. The five mutations operators are: (i) *Absolute value insertion (ABS)* in which each numerical expression (variable or method call) or literal is replaced by its absolute value. (ii) *Arithmetic operator replacement (AOR)* in which each arithmetic expression using Java arithmetic operators `+`, `-`, `*`, `/`, `%` is replaced by a new arithmetic expression with same operands but where the operator is changed into another one of the same family, chosen uniformly at random. Two other mutation candidates are also the left and the right operand alone, after removing the operator and one of the two operands; (iii) *Logical connector replacement (LCR)* in which each logical expression using Java logical operators `&&` and `||` is replaced by a new logical expression with the same operands, but where the logical operator is changed by another one. Moreover, each logical expression is also mutated by the constants `true` and `false`. Two other mutation candidates are also the left and the right operand alone, after removing the operator and one of the two operands. (iv) *Relational operator replacement (ROR)* in which each relational expression using Java relational operators `<`, `<=`, `>`, `>=`, `==` and `!=` is mutated to a relational expression with the same operands but where the relational operator is changed with another. Moreover, each relational expression is replaced by the constants `true` and `false`; (v) *Unary operator inversion (UOI)* where each arithmetic and logical expression are mutated. Arithmetic ones are mutated to their negative value (*i.e.* multiplied by `-1`), their incremented value (*i.e.* add `1`) and their decremented value (*i.e.* subtract `1`). Logical expression are complemented (*i.e.* apply the `not` (`!`) Java operator).

3.3 Results

In this section, we answer our research questions. All the empirical data that are behind the figures and interpretation are publicly available online[3].

Research Question 1 Do our learning algorithms improve the accuracy of impact prediction?

Table 2 gives the value of the evaluation metrics presented in Section 3.2.1. The first, second, and third columns are respectively the name of the package, the mutation operator and the number of mu-

[2] `http://cloc.sourceforge.net/`

[3] `https://github.com/v-m/ChangePropagation-dataset`

Table 2: Precision, recall and f-score obtained for change impact prediction with 3 techniques: transitive closure (TC), Algorithm Binary and Algorithm Dichotomic. The bold faced values indicate the best results for a single metric (up to rounding precision).

			Precision			Recall			F-score		
Package	Op.	#mut.	TC	Bin	Dic	TC	Bin	Dic	TC	Bin	Dic
Codec	ABS	600	0.15	0.28	**0.44**	0.93	**0.95**	0.89	0.17	0.30	**0.48**
	AOR	600	0.25	**0.41**	0.33	0.92	**0.93**	0.92	0.29	**0.45**	0.34
	LCR	600	0.15	0.26	**0.35**	**0.93**	**0.93**	0.90	0.20	0.29	**0.36**
	ROR	600	0.24	0.32	**0.36**	0.90	0.91	**0.92**	0.27	0.36	**0.38**
	UOI	600	0.27	0.40	**0.46**	**0.91**	**0.91**	0.85	0.30	0.43	**0.46**
	All	3000	0.21	0.33	**0.39**	0.92	**0.93**	0.90	0.25	0.37	**0.40**
Collections	ABS	600	0.55	0.75	**0.81**	0.80	0.80	**0.83**	0.43	0.59	**0.67**
	AOR	600	0.65	0.78	**0.81**	0.68	**0.70**	0.63	0.44	**0.55**	0.51
	LCR	600	0.56	0.69	**0.72**	0.77	**0.78**	0.77	0.43	0.51	**0.53**
	ROR	600	0.59	0.80	**0.84**	0.64	**0.68**	0.66	0.40	0.55	**0.56**
	UOI	600	0.63	0.80	**0.81**	0.68	**0.70**	**0.70**	0.44	**0.57**	0.56
	All	3000	0.60	0.76	**0.80**	0.71	**0.73**	0.72	0.43	0.55	**0.57**
IO	ABS	600	0.36	0.49	**0.54**	0.91	**0.92**	0.85	0.32	0.44	**0.46**
	AOR	600	0.41	0.63	**0.81**	0.93	**0.94**	0.90	0.41	0.63	**0.77**
	LCR	539	0.33	0.37	**0.61**	**0.92**	0.91	0.74	0.35	0.39	**0.46**
	ROR	600	0.44	0.50	**0.59**	0.90	**0.92**	0.89	0.44	0.50	**0.58**
	UOI	497	0.43	0.51	**0.69**	0.93	**0.94**	0.86	0.44	0.54	**0.65**
	All	2836	0.39	0.50	**0.65**	0.92	**0.93**	0.85	0.39	0.50	**0.58**
Lang	ABS	600	0.30	0.68	**0.69**	**0.99**	**0.99**	0.98	0.33	0.69	**0.70**
	AOR	600	0.47	0.65	**0.67**	**0.99**	**0.99**	**0.99**	0.52	0.68	**0.70**
	LCR	600	0.38	0.55	**0.59**	**0.98**	0.96	0.93	0.43	0.59	**0.62**
	ROR	600	0.49	0.66	**0.67**	**0.98**	0.97	0.97	0.54	0.68	**0.69**
	UOI	600	0.51	0.74	**0.75**	**0.98**	**0.98**	0.97	0.55	0.75	**0.77**
	All	3000	0.43	0.66	**0.67**	**0.98**	**0.98**	0.97	0.47	0.68	**0.70**
Shindig	ABS	600	0.34	0.49	**0.55**	0.95	**0.97**	**0.97**	0.33	0.49	**0.54**
	AOR	600	0.43	0.70	**0.79**	0.94	0.95	**0.96**	0.42	0.70	**0.78**
	LCR	600	0.49	0.48	**0.61**	0.88	**0.89**	0.87	0.45	0.43	**0.55**
	ROR	600	0.51	0.55	**0.63**	0.87	**0.88**	0.87	0.45	0.52	**0.57**
	UOI	530	0.48	0.66	**0.71**	0.91	**0.93**	**0.93**	0.44	0.62	**0.69**
	All	2930	0.45	0.58	**0.66**	0.91	**0.92**	**0.92**	0.42	0.55	**0.63**
Sonar	ABS	600	0.53	0.68	**0.69**	0.20	0.22	**0.25**	0.21	0.25	**0.28**
	AOR	538	0.39	0.84	**0.89**	0.27	**0.29**	0.28	0.22	**0.31**	0.30
	LCR	339	0.67	0.61	**0.82**	0.25	**0.33**	0.28	0.22	**0.33**	0.30
	ROR	600	0.62	0.65	**0.66**	0.35	**0.42**	**0.42**	0.32	0.42	**0.43**
	UOI	347	0.46	**0.86**	0.83	0.27	0.29	**0.30**	0.23	**0.31**	**0.31**
	All	2424	0.53	0.73	**0.78**	0.27	**0.31**	**0.31**	0.24	**0.32**	**0.32**
Spojo	ABS	48	0.11	0.83	**0.87**	**0.94**	**0.94**	0.92	0.06	0.78	**0.80**
	AOR	50	0.26	0.30	**1.00**	0.65	**0.68**	0.62	0.13	0.17	**0.62**
	LCR	59	0.22	0.32	**0.42**	**1.00**	**1.00**	0.97	0.32	0.42	**0.48**
	ROR	190	0.30	0.41	**0.52**	0.92	**0.94**	0.93	0.36	0.47	**0.54**
	UOI	37	0.29	0.78	**0.85**	0.92	0.93	**0.95**	0.25	0.74	**0.82**
	All	384	0.24	0.53	**0.73**	0.89	**0.90**	0.88	0.22	0.52	**0.65**
Total		17574	0.41	0.58	**0.67**	0.80	**0.81**	0.79	0.35	0.50	**0.55**

tants considered. Then we have three multi-columns, one for each metric (precision, recall and f-score). Each multi-column is made of three columns which are the value obtained using the transitive closure (TC) baseline, the value obtained with the binary algorithm, and the value obtained with the dichotomic algorithm. These values are the average of ten-fold cross validation. For each multi-column, the best value is shown in bold face.

First, if we compare the performance of the transitive closure algorithm (TC) and the two learning algorithms, we observe that in 100% of cases, both the precision and the f-score are improved using our learning algorithms. For the recall, in 25 cases out of 35 our learning algorithm gives better recall scores and in 7 cases out of 35 they are equivalent.

Figure 3 shows a scatter plot of the average F-scores obtained for all projects and mutation oper-

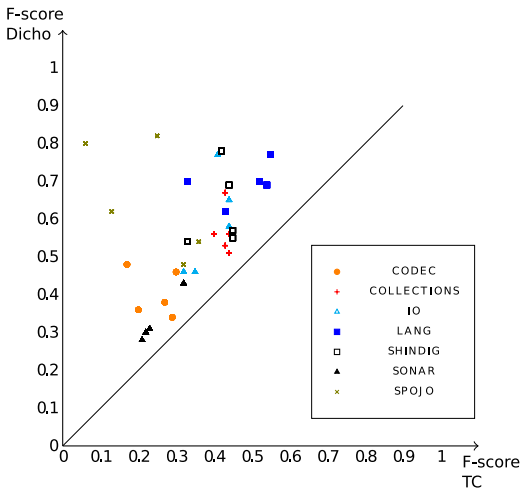

Figure 3: The performance improvement of the dichotomic algorithm over the basic TC impact prediction. One point represents a mutation operator for a given project, and is located at coordinates $(x = F_{TC}, y = F_{dichotomic})$.

ators for TC (x-axis) and dichotomic (y-axis). The line represents $y = x$. Since all points are in the upper-left part, above $y = x$, this figure graphically highlights that dichotomic much improves the prediction. In addition, we also graphically note that some operators are really far from $y = x$, which mean a strong improvement.

For the precision, the best improvement is for the Spojo project with the ABS operator where the precision raises from 0.11 for TC to 0.83 (+0.72) using binary algorithm and 0.87 (+0.76) using the dichotomic algorithm. Overall, the binary and the dichotomic algorithms provide respectively an average precision improvement of 0.18 and 0.26 when averaged over all projects and all mutation operators.

For the recall, the values are quite stable, the best improvement for binary and dichotomic algorithms are both with Sonar. For binary, the best improvement is obtained with the LCR operator for which the recall is improved from 0.25 for the baseline to 0.33 (+0.08). For dichotomic, the best improvement is obtained with the ROR operator for which the recall is improved from 0.35 for the baseline to 0.42 (+0.07). However, there is no real improvement with our learning algorithm compared to the transitive closure baseline on average (±0.01).

For the f-score, the best improvement is also obtained on the Spojo project with the ABS operator where the f-score raises from 0.06 using the transitive closure baseline to 0.78 (+0.72) using the binary algorithm, and 0.8 (+0.74) using the dichotomic approach. Globally, the binary and the dichotomic algorithms lead respectively an average f-score improvement of 0.15 and 0.20.

If we compare the two approaches, we can see that

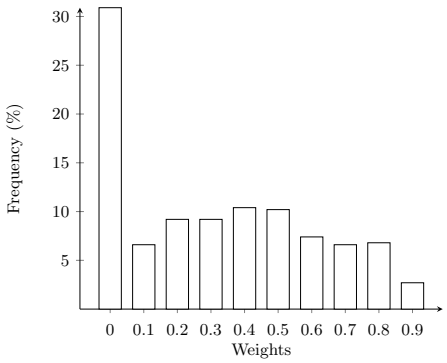

Figure 4: Histogram of the weights learned using the dichotomic algorithm aggregated over all projects of our dataset, all mutation operators, with a CHA based call graph. Each bin represents the amount of weights whose value lies in the range $[x, x + 0.1[$, where x is the value indicated under the bin.

the dichotomic algorithm has better precision values than the binary for 33 cases out of 35. But, dichotomic only has better recall in 6 cases out of 35. F-scores are also better for the dichotomic algorithm in 29 cases out of 35 (in 1 case they have the same score) than the binary one. All those observations suggest that dichotomic algorithm is better than binary.

To ensure that the performance difference between dichotomic and binary is significantly different, we run a Mann-Whitney-Wilcoxon statistical test on the precisions, recalls and f-scores. The null hypothesis (H_0) is: "the metrics (either precision, or recall, or f-score) obtained for the binary and the dichotomic algorithms are drawn from the same population". The obtained p-values are, for precision: 10×10^{-16}, for recall: 4.1×10^{-12} and for f-score: 8.7×10^{-14}. For each metric, the conclusion is then that the null hypothesis is rejected, which means that dichotomic is better in a statistically valid sense.

To sum up, mixing simple machine learning techniques with call graphs can achieve better performances for change impact prediction than a standard transitive closure approach. Among the two algorithms we propose in this paper, the best algorithm is the dichotomic approach.

Research Question 2 For the dichotomic approach, is the threshold value used for prediction important?

We now focus on dichotomic only. Figure 5 shows the f-scores obtained for each project and each operator, using different threshold values. Recall that a threshold t is used at prediction time, if the weight of an edge is higher than t, we consider that the impact propagates, otherwise not. Each sub-graph is a different project, each curve is a different operator.

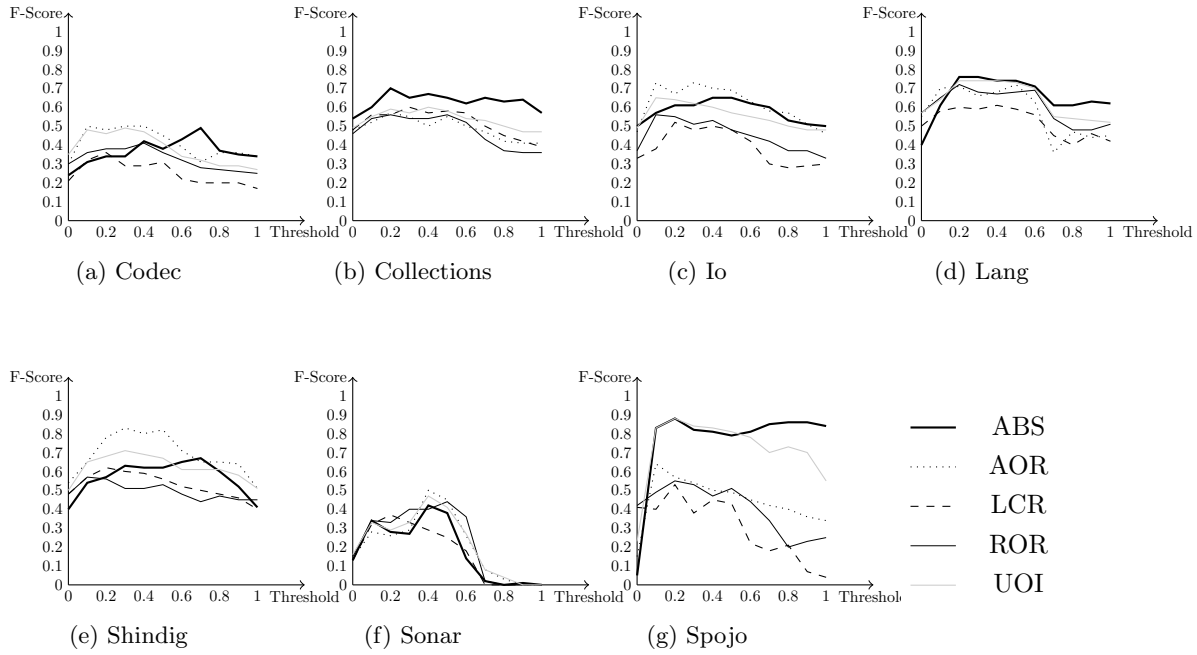

(a) Codec (b) Collections (c) Io (d) Lang

(e) Shindig (f) Sonar (g) Spojo

Figure 5: The impact of the prediction threshold on the F-Score. Low (< 0.1) and high (> 0.5) thresholds yield bad performances, otherwise any threshold between is virtually equivalent.

In general, we notice an inverse U shape for most curves. This means that a very low threshold yields bad performances, as well as high threshold. This can be explained as follows. Low thresholds means that most edges are said to propagate the impact. In this case, the precision drops too much [4]. On the other hand, high thresholds mean that very few edges are said to propagate. Since the dichotomic convergence slowly goes from 0 to the empirical probability, the bad performance of high threshold values is due to the fact the threshold is too high with respect to the convergence.

Overall, the best threshold value is 0.2, even if any other values between 0.1 and 0.6 give similar results, yet only slightly worse results.

Figure 4 shows an histogram of the weights obtained with the dichotomic algorithm over the 7 software projects. We observe that most edges have a weight between 0 and 0.1. Then the distribution is rather flat, exhibiting a local maximum around 0.4. This figure confirms that learning indeed happens. The edges that have never been visited are still at zero. The edges that are involved only in few propagations have a low weight. Consequently, a threshold around 0.2 indeed drops all edges with a low weight, those that are unlikely to propagate an impact.

To predict that an edge propagates an impact or not, very low and very high threshold are both bad. A good threshold value lies around 0.2. Beyond 0.6, the performance drops.

[4] indeed, the F-score for a threshold of zero is exactly the one of TC

Research Question 3 If we do not take into consideration the class hierarchy analysis in the call graph, what is the impact on the learning time and the prediction accuracy of dichotomic algorithm?

Table 3 gives the f-scores values for both binary and dichotomic approaches on call graph with and without class hierarchy analysis. The first and second columns are respectively the name of the package and the mutation operator. The third and fourth columns are respectively the median f-score for the binary approach without and with class hierarchy analysis. The fifth and sixth columns are respectively the median f-score for the dichotomic approach without and with class hierarchy analysis.

Let us first concentrate on the results of the binary approach. We observe that we have a better f-score in 6 cases out of 35 with class hierarchy analysis versus 25 without (and 4 cases for which the f-score is the same). Thus, using a call graph without class hierarchy analysis gives better results with the binary algorithm. The most representative projects are Apache Common Collections with AOR mutation operator (0.55 with class hierarchy analysis versus 0.64 without) and Apache Commons IO with ROR operator (0.5 with class hierarchy analysis versus 0.59 without). In both cases, the f-score is improved of 0.09 by just removing inheritance and interface considerations from the call graph. This is logical: a CHA call graph adds many edges: since binary switches edges from "not propagating" (weight of zero) to "propagating" (weight of one), the additional edges are all switched far too often. To some extent, when we con-

Table 3: The F-score obtained for binary and dichotomic algorithms when applied to a call graph with and without class hierarchy analysis. The results are averaged over 10-fold cross validation.

		Binary		Dichotomic	
Package	Op.	-CHA	+CHA	-CHA	+CHA
Codec	ABS	0.36	0.30	0.40	0.48
	AOR	0.47	0.45	0.49	0.34
	LCR	0.31	0.29	0.33	0.36
	ROR	0.42	0.36	0.41	0.38
	UOI	0.45	0.43	0.48	0.46
Collections	ABS	0.63	0.59	0.70	0.67
	AOR	0.64	0.55	0.67	0.51
	LCR	0.53	0.51	0.57	0.53
	ROR	0.57	0.55	0.55	0.56
	UOI	0.61	0.57	0.63	0.56
IO	ABS	0.45	0.44	0.47	0.46
	AOR	0.63	0.63	0.76	0.77
	LCR	0.45	0.39	0.53	0.46
	ROR	0.59	0.50	0.58	0.58
	UOI	0.56	0.54	0.66	0.65
Lang	ABS	0.70	0.69	0.72	0.70
	AOR	0.73	0.68	0.72	0.70
	LCR	0.60	0.59	0.61	0.62
	ROR	0.71	0.68	0.70	0.69
	UOI	0.74	0.75	0.73	0.77
Shindig	ABS	0.54	0.49	0.55	0.54
	AOR	0.70	0.70	0.74	0.78
	LCR	0.43	0.43	0.59	0.55
	ROR	0.52	0.52	0.53	0.57
	UOI	0.63	0.62	0.71	0.69
Sonar	ABS	0.28	0.25	0.27	0.28
	AOR	0.30	0.31	0.29	0.30
	LCR	0.35	0.33	0.30	0.30
	RO R	0.39	0.42	0.36	0.43
	UOI	0.26	0.31	0.27	0.31
Spojo	ABS	0.82	0.78	0.77	0.80
	AOR	0.19	0.17	0.54	0.62
	LCR	0.39	0.42	0.49	0.48
	ROR	0.44	0.47	0.51	0.54
	UOI	0.77	0.74	0.78	0.82

sider a graph without CHA, the decrease of recall due to losing some edges is compensated in the F-Score by the increase of precision in the other cases.

For the dichotomic approach, we observe that we have a better f-score in 16 cases out of 35 with class hierarchy analysis with regards to 17 without (and 2 cases for which the f-score is the same). According to the F-Score, taking or not the class hierarchy into consideration is equivalent.

To strenghten our observations, we run two Mann-Whitney statistical tests, one for binary algorithm and one for dichotomic one. On each test, the null hypothesis (H_0) is: "the f-scores obtained for the call graph with and without class hierarchy analysis are drawn from the same population". For the binary approach, we obtain a p-value of 1.9×10^{-4}. The null hypothesis is rejected at significance level of 0.05, which confirms our observations. For the dichotomic approach, we obtain a p-value of 0.1124. Here, the null hypothesis cannot be rejected at significance level of 0.05 which means we cannot conclude anything statistically. In other words, they may be drawn from the same population, which comfort us in our observations regarding the fact that considering CHA or not is equivalent in this setup.

Table 4: Execution time of learning propagation information for binary and dichotomic algorithms. The values are averaged over all mutation operators.

	Binary		Dichotomic	
Package	-CHA	+CHA	-CHA	+CHA
Codec	1s	2s	2s	2s
Collections	2s	1m2s	3s	2m13s
IO	1s	1s	1s	10s
Lang	21s	35s	8s	32s
Shindig	1s	3s	370ms	1s
SonarQube	500ms	600ms	288ms	662ms
Spojo	200ms	250ms	300ms	267ms

Now we consider the execution time. Table 4 reports the mean execution time during learning for the binary approach and the dichotomic approach (over all mutation operators). The first column is the package name, the second and third columns are the execution time required to run the binary algorithm with and without class hierarchy analysis and the fourth and fifth report the same values but for the dichotomic. We observe that excluding class hierarchy analysis yields a 1x to 44x improvement. The main explanation is that learning time is correlated to the number of edges. The correlation is not linear since we compute all paths between all pairs of application nodes and test nodes; and the path computation is not linear.

We see that the increase in learning time with CHA phenomenon is project dependent: for Apache Commons Codec the difference in learning time is quite low (around 1 to 2 times faster) but for some others, such as Apache Commons Collections the difference is more important (31 or 44 depending on the learning algorithm). This confirms our explanation based on the number of edges: Apache Common Collections extensively uses inheritance. Table 1 shows that including the class hierarchy analysis in our call graph adds 2795 new edges to the call graph of Apache Common Collections. This is the greatest increase of number of edges, to which corresponds the greatest increase in computation time as well.

Indeed, we have noticed that learning time becomes untractable for very large CHA call graphs extracted from large software projects. By disabling CHA, we were able to successfully run our learning algorithm on three larger projects: Google Gson, Jodatime and JGit which have respectively 20072, 83124 and 133856 lines of codes. This preliminary result paves the way for future experiments on learning for impact predication on large software.

To sum up, in our experiments, using a call graph without class hierarchy does not result in a noticeable performance decrease for Algorithm

dichotomic one. However, it much speeds-up the learning process depending on the project, which enables to analyze larger projects.

3.4 Threats to Validity

Our results are of computational nature. A major bug in our software can invalidate our findings. We have published all our code on Github so to facilitate reproduction and falsification of our results, if necessary.

In this paper, we use 600 mutants per mutation operator and per project. Using ten-fold cross validation, this results in testing sets of up to 60 items. Since the aggregate performance measures (precision, recall and f-score) are rather stable over folds, we have confidence that this is enough to back up our conclusions. However future work with more mutants is required.

The impact prediction much depends on the structure of the call graphs [14]. For instance, the presence of large utility methods, with many incoming and outgoing edges has a direct impact on the predication performance. Our results may only be valid for Java software, or even worse only valid for the projects under study. Future work in this field will strengthen the external validity of our findings.

4 Related work

4.1 Machine Learning and Software Engineering

Machine Learning has already been applied by several authors in software engineering. All those applications want to solve a great amount of different problems. Srinivasan and Fisher have applied decision trees and neural networks for estimating development effort [15]. Elish and Elish used support vector machines to predict defect-prone modules in four projects [16]. Tian and Noore have predicted software cumulative failure time using neural networks [17]. For bug triage, Jeong and Kim have used machine learning on bug tossing graphs (graph exposing the reassignment of a bug) [4] and Kim et al. have used machine learning on crash graphs (aggregated view of multiple crash) [18]. Baskiotis et al. [19] used learning on control flow graph to generate new paths for testing. All of those works use machine learning approaches but none of them have applied machine learning to change impact analysis.

4.2 Mutation Testing

Mutation testing is an old concept which has seen many contributions over years. Jia and Harman propose a survey regarding this topic [20]. We focus here on on the works that are most related to ours.

Strug and Strug [21] use data and control flow graphs and classification for detecting similar mutants. They use similar tools as us (learning and mutation) to reduce the number of mutant considered when doing mutation testing. In our side, we use those tools for change impact analysis.

Do and Rothermel [22] described a protocol to study test case prioritization techniques based on mutation. Their protocol and ours share the same intuition, use test cases to determine which tests is impacted by the change. However, we have a different goal, they study test case prioritization, we study impact prediction.

4.3 Change Impact Analysis

Change impact analysis has been studied for many years and many algorithms have been proposed. Many categorizations of such algorithms exists. Bohner and Arnold proposed two types of analysis: dependency analysis and traceability analysis [1]. The former analyzes the source-code of the program in a relatively fine granularity (*e.g.* methods call, data usage, control statements,...) while the latter compares elements from coarser granularity such as documentation and specifications (*e.g.* UML, *etc.*). Moreover, different types of impact-determination techniques are presented. According to this paper, our approach is a dependency analysis based on a transitive closure technique.

Bohner and Arnold [1] and Li et al. [2] list the notable graph based approaches. Different types or variants of software graphs have been used to perform change impact analysis, a common example is the program dependence graphs (*a.k.a.* PDG) [5]. In this paper we focus on call graph.

Walker et al. [23] have proposed an impact analysis tool named TRE. Their approach uses conditional probability dependency graphs where nodes are classes and edges are added when a class A depends on a class B if A contains anything resolving to B. The conditional probabilities are extracted from the CVS repository and defined by the number of time two classes are changed on the same commit. Then technical risk for a change is determined based on the CP graph. They work on class granularities and give no concrete information about the evaluation. In contrast to this work, we work on a finer granularity (methods) which gives us more realistic data and give numerical evidences for 7 Java packages.

Hattori et al. [11] have used an approach based on call graphs to study the propagation. Their evaluation is made on a small dataset made of three projects. Their goal was to show that the precision and recall are good tools as evaluation of the performance for an impact analysis technique. In contrast,

we propose a concrete impact analysis technique.

Zimmermann and Nagappan [24] have proposed to use dependency graphs to estimate the most critical parts of a software. Their approach uses network measures and complexity metrics to make the predictions. They assess their findings based on one popular but proprietary software, where they are able to determine part of the software that can cause issues. In contrast, we propose a technique to determine which part of a software will be impacted by a potential change. Moreover, we experiment our approach on seven different open-source software packages.

Law and Rothermel [25] have proposed an approach for impact analysis; their technique is based on a code instrumentation to analyze execution stack traces. They compare their technique against simple call-graphs on one small software. Our technique is compared to similar graphs but with seven different software packages.

Antoniol et al. [26] also address impact analysis. However, they consider a slightly different problem setting, because they take as input a bug report or a modification request and not a single source code element as we do. The same argument applies for the recent work by Gethers and colleagues [27].

In conclusion, none of them has proposed to use machine learning to improve their technique.

5 Conclusion

In this paper we have explored the possibilities of using learning to improve impact prediction based on call graphs. We have shown that we can improve the prediction performance with a learning algorithm based on the addition of weights to the call graph edges.

From a software engineering perspective, our results show that having a learning perspective on change impact prediction is a promising research direction. From a machine learning perspective, the algorithms introduced in this paper can be considered as basic. More machine learning work has to be done to further study this new problem of learning impact weights in a graph.

References

[1] S. A. Bohner and R. S. Arnold, *Software Change Impact Analysis*. Los Alamitos, CA, USA: IEEE Computer Society Press, 1996.

[2] B. Li, X. Sun, H. Leung, and S. Zhang, "A Survey of Code-based Change Impact Analysis Techniques," *Software Testing, Verification and Reliability*, vol. 23, no. 8, pp. 613–646, Dec. 2013.

[3] R. J. Turver and M. Munro, "An Early Impact Analysis Technique for Software Maintenance," *Journal of Software Maintenance: Research and Practice*, vol. 6, no. 1, pp. 35–52, Jan. 1994.

[4] G. Jeong, S. Kim, and T. Zimmermann, "Improving Bug Triage with Bug Tossing Graphs," in *Proceedings of the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering*, ser. ESEC/FSE '09. New York, NY, USA: ACM, 2009, pp. 111–120.

[5] J. P. Loyall and S. A. Mathisen, "Using Dependence Analysis to Support the Software Maintenance Process," in *Proceedings of the Conference on Software Maintenance*, ser. ICSM '93. Washington, DC, USA: IEEE Computer Society, 1993, pp. 282–291.

[6] S. Bohner, "Software Change Impacts - An Evolving Perspective," in *Proceedings of the International Conference on Software Maintenance*, ser. ICSM '02, 2002, pp. 263–272.

[7] R. S. Arnold and S. A. Bohner, "Impact Analysis - Towards a Framework for Comparison," in *Proceedings of the Conference on Software Maintenance*, ser. ICSM '93. Washington, DC, USA: IEEE Computer Society, 1993, pp. 292–301.

[8] T. A. Budd, R. J. Lipton, R. A. DeMillo, and F. G. Sayward, "Mutation Analysis," Yale University, Department of Computer Science, Tech. Rep., 1979.

[9] D. Grove, G. DeFouw, J. Dean, and C. Chambers, "Call Graph Construction in Object-oriented Languages," in *Proceedings of the 12th Conference on Object-oriented Programming, Systems, Languages, and Applications*, ser. OOPSLA '97. New York, NY, USA: ACM, 1997, pp. 108–124.

[10] R. Kohavi, "A Study of Cross-validation and Bootstrap for Accuracy Estimation and Model Selection," *International joint conference on artificial intelligence, 1995*, pp. 1137–1143, 1995.

[11] L. Hattori, D. Guerrero, J. Figueiredo, J. Brunet, and J. Damásio, "On the Precision and Accuracy of Impact Analysis Techniques," in *Proceedings of the Seventh IEEE/ACIS International Conference on Computer and Information Science (Icis 2008)*, ser. ICIS '08. Washington, DC, USA: IEEE Computer Society, 2008, pp. 513–518.

[12] K. N. King and A. J. Offutt, "A Fortran Language System for Mutation-based Software Testing," *Software: Practice and Experience*, vol. 21, no. 7, pp. 685–718, Jul. 1991.

[13] A. J. Offutt, A. Lee, G. Rothermel, R. H. Untch, and C. Zapf, "An Experimental Determination of Sufficient Mutant Operators," *ACM Transactions on Software Engineering and Methodology*, vol. 5, no. 2, pp. 99–118, Apr. 1996.

[14] V. Musco, M. Monperrus, and P. Preux, "A Generative Model of Software Dependency Graphs to Better Understand Software Evolution," *ArXiv e-prints, 1410.7921v2*, Oct. 2014.

[15] K. Srinivasan and D. Fisher, "Machine Learning Approaches to Estimating Software Development Effort," *IEEE Transactions on Software Engineering*, vol. 21, no. 2, pp. 126–137, Feb. 1995.

[16] K. O. Elish and M. O. Elish, "Predicting Defect-prone Software Modules using Support Vector Machines," *Journal of Systems and Software*, vol. 81, no. 5, pp. 649–660, May 2008.

[17] L. Tian and A. Noore, "Evolutionary Neural Network Modeling for Software Cumulative Failure Time Prediction," *Reliability Engineering & System Safety*, vol. 87, no. 1, pp. 45–51, Jan. 2005.

[18] S. Kim, T. Zimmermann, and N. Nagappan, "Crash Graphs: an Aggregated View of Multiple Crashes to Improve Crash Triage," in *Proceedings of the 41st International Conference on Dependable Systems Networks*, ser. DSN '11, Jun. 2011, pp. 486–493.

[19] N. Baskiotis, M. Sebag, M.-C. Gaudel, and S.-D. Gouraud, "Exist: Exploitation/Exploration Inference for Statistical Software Testing," in *On-line Trading of Exploration and Exploitation, NIPS 2006 Workshop*, 2006.

[20] Y. Jia and M. Harman, "An Analysis and Survey of the Development of Mutation Testing," *IEEE Transactions on Software Engineering*, vol. 37, no. 5, pp. 649–678, Sep. 2011.

[21] J. Strug and B. Strug, "Machine Learning Approach in Mutation Testing," in *Testing Software and Systems*, ser. Lecture Notes in Computer Science, B. Nielsen and C. Weise, Eds. Springer Berlin Heidelberg, 2012, no. 7641, pp. 200–214.

[22] H. Do and G. Rothermel, "A Controlled Experiment Assessing Test Case Prioritization Techniques via Mutation Faults," in *Proceedings of the 21st International Conference on Software Maintenance*, ser. ICSM '05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 411–420.

[23] R. J. Walker, R. Holmes, I. Hedgeland, P. Kapur, and A. Smith, "A Lightweight Approach to Technical Risk Estimation via Probabilistic Impact Analysis," in *Proceedings of the International Workshop on Mining Software Repositories*, ser. MSR '06. New York, NY, USA: ACM, 2006, pp. 98–104.

[24] T. Zimmermann and N. Nagappan, "Predicting Defects Using Network Analysis on Dependency Graphs," in *Proceedings of the 30th International Conference on Software Engineering*, ser. ICSE '08. New York, NY, USA: ACM, 2008, pp. 531–540.

[25] J. Law and G. Rothermel, "Whole Program Path-Based Dynamic Impact Analysis," in *Proceedings of the 25th International Conference on Software Engineering*, ser. ICSE '03. Washington, DC, USA: IEEE Computer Society, 2003, pp. 308–318.

[26] G. Antoniol, G. Canfora, G. Casazza, and A. de Lucia, "Identifying the Starting Impact Set of a Maintenance Request: A Case Study," in *Proceedings of the Conference on Software Maintenance and Reengineering*, ser. CSMR '00. Washington, DC, USA: IEEE Computer Society, 2000, pp. 227–.

[27] M. Gethers, B. Dit, H. Kagdi, and D. Poshyvanyk, "Integrated Impact Analysis for Managing Software Changes," in *Proceedings of the 34th International Conference on Software Engineering*, ser. ICSE '12. Piscataway, NJ, USA: IEEE Press, 2012, pp. 430–440.